

PRIVACY COMPLIANT MULTIPLE DATASET CORRELATION SYSTEM

INVENTORS:

G. Colby Conkwright
7715 13th Avenue NW
Bradenton, FL 34209
Citizen of: USA

Michael J. Vinson
5017 Sturbridge Court
Sarasota, FL 34238
Citizen of: USA

Frank S. Foster IV
5905 Cherry Oak Drive
Valrico, FL 33594
Citizen of USA

ASSIGNEE:

erinMedia, Inc.
2509 63rd Avenue East
Bradenton, FL 34203

ATTORNEY:

Greenberg Traurig
1750 Tysons Boulevard, 12th Floor
McLean, VA 22102
(703) 749-1377

REFERENCE TO RELATED APPLICATIONS

The present application claims priority from Provisional U.S. Patent Application Serial Number 60/176,177, filed January 13, 2000, and the Provisional U.S. Patent Application is incorporated by reference in its entirety.

5 FIELD OF THE INVENTION

The present invention relates to the fields of data collection and data analysis. In particular, the present invention provides a system and method for privacy-secure data collection and correlation of such data with data from other sources.

BACKGROUND OF THE INVENTION

10 Advertisers tend to group prospective customers into broad demographic and geographic categories, possibly due to limitations in currently available market research methods with respect to determination of the effect of their advertisements. In addition, they use information gleaned from data mining to mass-market products to groups of prospective buyers. Unfortunately, the data searched during this data mining often
15 contains low-validity information that is derived from small sample populations.

Due to these inherent data validity problems, statistics generated by such data mining may not accurately reflect a given market. That is, the statistics may not mean that all persons in a group will buy a product, but rather they imply that some person in a group may have a higher probability of buying the product than someone in another
20 categorized group. For example, the data mining may show that more scuba equipment could be sold to 20-40 year-olds in Miami than to 50-80 year-olds in Kansas City.

Based on this data, advertisers carefully select the television shows, magazines, billboards, or other media on or in which their advertisements run. In the case of television, advertisers traditionally gravitate toward programs that garner higher ratings
25 for desired audiences and then select advertising slots within those shows. Advertisers purchase ratings data from market research organizations, who collect and analyze data on the viewing habits of individuals and then publish the results.

Examples of such research organizations include A. C. Nielsen and Arbitron. Such companies typically monitor television-viewing habits of a relatively small number

of viewers through telephone polls, specialized set-top monitoring "Nielsen" boxes, or viewer diaries. The results of these surveys are then extrapolated to the population at large.

As can be expected, extrapolation of small-population data to the population at large is prone to many different limitations, with accuracy perhaps the most notable. For example, if there were only 200 persons over 65 years of age in a sample, their compiled viewing behaviors may be purported to be representative of the viewing behaviors of the 35 million people in the U.S. over 65 years of age.

Obviously, larger, more random sample populations are preferred over smaller sample populations. This is true because a larger sample population tends to reduce the impact of suspect behaviors. Such suspect behavior might include distorted or inaccurate information provided in written television viewing logs, or intentionally leaving the television "on" to a certain channel to insure higher ratings for a desired show even if the individual being sampled is not watching that show. If the behavior of even one of the 200 persons in the previous example was suspect, this may translate to errors in the predictions of approximately 175,000 people; if the sample population is increased to 50,000 people, an individual whose behavior was suspect would translate into prediction errors for only approximately 700 people. As advertisers continue to base their decisions on small-sample data, they are continuing to question whether their advertisements are reaching intended audiences.

While accuracy is certainly a big problem in the prior art, it is not the only problem. Another limitation is the specificity with which behaviors may be inferred as they pertain to specific demographic groups. For example, if only one of 200 sampled senior citizens is a single Asian with no dependents and has an annual income over \$100,000, making an inference based on this more specific group is likely to be highly inaccurate; in many cases the behaviors of an entire demographic sub-group are attributed to the sampled behavior of only one person.

Another factor contributing to the inaccuracy of prior art is reliability. Invasive sampling methods such as those described above can cause many problems, including determining how much of the data can be trusted. Sampled individuals may not be

willing to disclose, for example, that they watch adult (e.g., X-rated) programming or other controversial programming. Without such information, all data generated becomes unreliable.

Still another problem is that even if the sample data can be trusted, the memory of
5 a sampled individual or the ability of a sampled individual to adhere to documented
guidelines may not be accurate or complete. If a given individual is asked what they
watched last week, the likelihood that the response may be correct and specific is likely
to be low. Often, low response rates or missing journal information are extrapolated
according to previously collected data and rules determined therefrom. However, this
10 extrapolation is built on data generated through the inherently faulty means described
above.

The invasive sampling techniques used in the prior art also suffer from an
inherent flaw. Since these methods are invasive and participation is optional, differences
between the types of persons who may be willing to be sampled and those that are not
15 willing to be sampled may not be accounted for in such techniques.

While the effects of some of the problems in the prior art can be limited by
increasing the population sample size, population sample size increases are typically cost
prohibitive. The increased costs are the result of several factors, including equipment
purchase, installation, and repair; data collection and validation; and participant
20 compensation.

However, even when equipment and other costs are taken out of consideration
and larger samples are collected, such an increase in sampling size does not solve all of
the problems in the prior art. For example, the prior art also faces a problem with data
resolution. Most major media research organizations consider data in an all-or-nothing
25 fashion. For example, if a set of channels was watched during some sampling interval,
only the channel that was watched the most, or the one watched at the time of the sample,
would be counted, and it would be recorded as having been watched for the entire
sampling interval (typically anywhere from 30 seconds to several hours). Although some
in the prior art have attempted to mitigate this effect by sampling more frequently, there
30 is always the possibility that changes occurring between samples will be missed. Thus,

the use of data collection methods employed by the prior art tends to result in the generation of misleading or inaccurate viewing data.

5 Data collected by media research organizations and inferences resulting therefrom face still another problem; one of substance. The fact that overlapping data is collected across different medium types (digital, written, verbal, etc.) makes the determination of common denominators difficult, and thus renders objective statistical mining impossible. Inferences drawn from such data may only be lateral in nature, and cannot be readily mined for trends. For example, while the data collected may support the conclusion that one show is more popular than another, the particular reason why one is more popular
10 than the other cannot be extracted from this data. Such methods may be barely capable of supporting the most general popularity-type conclusions; any further analysis upon relationships of the conclusions is likely to be questionable at best, and accuracy may be lost each time more complicated, or deeper, inferences are drawn.

15 Unfortunately, there are many other problems with existing market research methodologies, such as the use of "Sweeps" or ratings periods, but most of these problems are at least partially statistically-correctable. However, the five major issues discussed - accuracy, group specificity, reliability, resolution, and data substance - are inherent to actively monitoring data within small samples and cannot be overcome by the prior art.

20 **SUMMARY OF THE INVENTION**

The data collection techniques used in the prior art arise from a model developed in the 1960's and 1970's. At that time, market research data collection and data transmission costs were very high, and a system of periodic sampling was established. A thirty second sampling window was chosen because a given household had an average of
25 three channels available, and "surfing" was a non-existent phenomenon; thus, it was a safe assumption that the same channel was watched for the entire sampling period. Today, most television viewers have over 65 channels available to them, and they are barraged with more commercials per hour on each channel. This combination gives viewers incentive to frequently change channels within what would be a sampling period
30 in the prior art.

Obviously, sampling methods from the 1970's are not capable of accurately representing television viewing habits in the year 2000 and beyond. Thus, a need exists to more accurately sample viewer data, insure the data collected is not suspect, infer from such data relationships and trends in viewing habits, project sampled data more
5 accurately to the population at large, and determine not only what shows, advertisements, or other content were watched, but also what portions of such content were watched. There also exists a need for ratings systems which can more accurately and objectively provide ratings of future programs. The data collection and data analysis aspects of the present invention can readily fulfill these needs.

10 A preferred embodiment of the present invention can provide advertisers with accurate ratings predictions of commercials and programs for specific demographic groups, rather than just providing overall ratings of programs which have already aired. While a preferred embodiment of the present invention involves television viewership statistics, the present invention can draw correlations between any dataset combinations,
15 such as, but not limited to, television program or commercial viewership and sales figures, or sales figures and demographics. The present invention may provide advertisers with a better understanding of both consumer needs and their own advertising needs.

One aspect of a preferred embodiment of the present invention provides a system
20 and method by which viewing behaviors of television viewers can be extracted electronically. While instrumentation and infrastructure development costs may be initially high, the present invention can allow data collection from a vast number of households without significant data collection, data storage, and data analysis cost increases as the number of households increases, or as the number of times data is
25 collected per household is increased.

The present invention takes a different approach to television market data collection than the prior art. Rather than periodically sampling user behavior, the present invention tracks user behaviors by recording set-top box events. Such a set-top box may record events including, but not limited to, set-top box state changes, such as a set-top
30 box being turned on or off, channel changes, volume changes, the use of an SAP feature,

or muting of particular content; the use of interactive content guides; Internet web site usage; and combinations of such events. Recorded set-top box events data may be periodically transmitted to a central data collection point where data analysis may begin, or such transmission may occur instantaneously. In a preferred embodiment, this data 5 collection method allows data to be gathered without requiring subjects to keep journals, push buttons, or even know their behaviors are being observed. This can be seen as an improvement over the prior art, as the invasive data collection methods used therein are likely to destroy data integrity.

The present invention also includes a method for mining compatible datasets so 10 that correlations and trends within and between the datasets can be uncovered. The present invention is tailored to the analysis of datasets that are extremely large; result from passive, privacy-secure data collection; and are relatively unbiased, such as datasets collected by set-top boxes described above. While the analysis of television marketing data is presently preferred, it will be apparent to one skilled in the art that the system and 15 method herein can be employed in other data collection and data analysis scenarios. Other contemplated embodiments include, but are not limited to, privacy-secure actuarial analysis, radio and Internet market data collection, and thought and behavioral predictions for artificial intelligence efforts and governmental planning.

The data mining and prediction portions of the present invention attempt to 20 uncover the interrelated rules that cause various data to arise, and a preferred embodiment does so while still respecting privacy concerns. Privacy can be maintained through anonymous data collection, which can be accomplished through a software upgrade to a standard set-top box. In a preferred embodiment, a satellite, cable, or other television provider (“cable company”) can provide a viewer with a set-top box which may be 25 specially instrumented to allow monitoring, recording, and transmission of set-top box events, as described above.

Traditional set-top boxes include a unique identification number (“ID”), and this 30 number can be used by a preferred embodiment of the present invention for identification purposes in lieu of personal information. To facilitate data analysis, a cable company can provide to the present invention a geographically associated code, such as, but not limited

to, a zip code or telephone number prefix, that corresponds with each set-top box. Given this information, set-top box ID's to be monitored can be chosen through various means, including, but not limited to, the present invention selecting set-top box ID's at random, the present invention selecting set-top box ID's based on geographic coverage, a cable 5 company selecting ID's based on its own criteria, or selecting all set-top box ID's. A combination of set-top box ID and geographically associated codes allows the present invention to maintain participant privacy while still allowing for determination of detailed demographic information through the inverse mathematical methods described herein.

10 Although privacy is an important part of the present invention, an alternative embodiment would allow set-top box operators to request a list of their viewing habits. This might be useful for parents or businesses wishing to monitor programs watched by their children or employees during a given day, or parents or businesses wishing to monitor other DATA1 datasets, such as internet viewing behaviors exhibited by 15 employees or family members.

Analysis of data collected through a privacy-oriented approach such as the set-top box method described above is inherently self-limiting, as only viewership information for a particular show or commercial can be determined for a given time over the sample population. While this may be of interest to advertisers, an advertiser's real concern is 20 that a show is reaching a particular target-market, and thus that they are spending their advertising money on shows which users of their products will watch. Thus, advertisers prefer detailed, grouped market research data, such as the ages, incomes, and other demographic information associated with a show's viewers. Through its novel data-correlation scheme, the preferred embodiment of the present invention can determine 25 such information while still maintaining the anonymity of those being sampled.

However, the present invention is not limited to providing only correlations between user behavior and demographic data; the present invention can draw correlations within and between any number of data sets with a common feature, such as the zip code of a given television viewer and the demographics associated with that zip code, regardless of the 30 data represented therein.

In a preferred embodiment, the present invention augments television viewing behavior data collected from set-top boxes with relatively static data from outside sources, such as, but not limited to, demographic information from demographic providers, news information from news providers, weather information from weather providers, and sales information from advertisers, manufacturers and producers. This information can be used not only to increase the number of categories into which individuals may be grouped, but also to take into account specific confounding events, such as a severe weather alert, a national or regional news story, a local school's play-off game, and special promotional offers. Demographic, regional, and other such relatively static data may be updated at intervals specific to the type of data collected.

The present invention may make certain assumptions based on collected data to reduce data storage requirements. These assumptions are well-fitted with the use of matrix manipulation schemes. For example, the present invention may assume individual age is intrinsic to a person, household income is intrinsic to a household, and weather patterns are intrinsic to geographic regions. Thus, demographic information comprising a matrix needs not be mutually exclusive. That is, if weather is a considered factor, weather data need not be collected and stored for every person in a geographic region, but simply can be held in one matrix that can be accessed for all people living within that geographic region.

The present invention can draw inferences within and between data sets through a variety of means. In a preferred embodiment, the present invention may use inverse mathematical methods to perform the desired analyses. These methods can be more simply expressed as techniques of linear and matrix algebra.

With established data mining and data comparison methods in place, another step is to extrapolate any calculations to the rest of the country. Through the data-collection methods described above, present invention can track viewing behaviors, demographic characteristics, and preferences associated with a geographically based group of people. If this same demographic information is obtained for the country as a whole, the present invention can use the data stored therein to extrapolate out its results to any region of the country, or to the country as a whole. The present invention includes, but is not limited

to, the development of an extrapolation system, which itself involves the use of mined trends and will evolve and improve over time. For example, "white males age 70-80" in the non-sample population may be considered to have more similarities with all persons age 70-80 in the sample than with whites alone due to a dominance of the age factor in whites. Due to characteristics particular to a given geographic location, such as, but not limited to, such a location including a resort community, or a local sports team in a playoff game, some factors, such as age, race, and gender, or even entire geographic regions, may be ignored in extrapolation procedures.

While some in the prior art have attempted to provide statistics similar to those available through the present invention, none have done so at confidence levels approaching those provided by the present invention. In addition, the present invention improves over the prior art by allowing the extrapolation of very specific cases, even when no data exists for that specific case. For example, if no Single Asian Males, 23-24 years old, with partial custody of 1 child, one previous marriage, with a B.S. in Chemistry, working as an assistant in a Chemical Laboratory, having an income between \$24,000 - \$27,000 per year, and living in a specific zip code in Miami, Florida were in the sample, the present invention could still calculate an anticipated behavior based upon combinations of subsets of these characteristics and their observed influences.

In fact, not only can the present invention infer viewing preferences and other behaviors for previously aired content, but, as additional sample data is collected, the present invention can also predict reactions to future content. The present invention can characterize previously aired content, such as a television program or a commercial, based on specific attributes thereof, such as volume changes; color changes; changes in brightness or contrast; speed of motion; background music mood; content; genres; actors appearing on the screen; plotlines; languages spoken; use of foul or offensive language; and the like. This information can then be cross-referenced against viewer reaction to that content, and suggestions can be made to make the content more appealing to a particular audience. With a database of viewer reactions to previously aired content, the present invention can also be used to analyze proposed content before it is aired, or even to suggest optimal programming content structure and substance.

It should be obvious to one skilled in the art that the system and method described in this specification are not constrained by the same limitations as traditional data collection and analysis techniques. The present invention provides non-invasive sample data collection, significantly increasing reliability. The viewing habits of increasingly 5 specific demographic groups can be ascertained while still maintaining high accuracy levels.

Additionally, the behavior data resolution is so fine as to allow the redefinition of television viewing behavior. For example, the exact percentage of each program that each group watched can be determined. In addition, the answer to "What percentage of 10 Group A watched at least 80% of Program 1 who watched less than 10% of Program 2 three weeks ago?" is just as easy for the computers to determine as a seemingly simpler question. Furthermore, event data for every set-top box and geographic region can be archived as a large, unbiased database of 1's and 0's; therefore mining the data for trends would entail literally 0% loss of related-accuracy. This is to say that the data's substance 15 is the same regardless of specificity of analysis.

It should also be obvious to one skilled in the art that the system and method described above can be used not only to rate television programs, but, unlike the prior art, the present invention can also rate television advertisements. The prior art is limited to television program ratings because the sampling periods required to accurately rate 20 television advertisements would result in more data than can be accurately collected, handled and characterized by the prior art. The novel data collection method described above provides viewer behavior data at a finer resolution than is possible through the sampling methods implemented by the prior art, and can thus be used to determine viewer behavior at any instant, including sections during a commercial.

Once data is collected, a further aspect of the present invention provides 25 advertisers or others wishing to analyze the data with an interactive interface for such data analysis. Such an interface can allow data analysis requests to be entered through a variety of interfaces, such as through command-line queries, graphical interfaces, or even as natural-language questions. Further, an output representation may be selected, 30 including, but not limited to, raw data, pie chart, time-progression, and the like. The

present invention may further track frequently requested analyses and automatically update those on a periodic basis to expedite delivery of such information.

While the present invention improves over the prior art when evaluated using current advertising schemes, the present invention can also allow a new type of advertising. Rather than an advertiser purchasing time during a given show that is broadcast to a large audience, many of whom may not be in a product's target audience, the present invention may allow advertisements to be delivered to only those set-top boxes whose viewers exhibit certain behaviors or exhibit a propensity toward specific products or services. This allows advertisers to directly reach those viewers who would be interested in an advertiser's product or service, thus decreasing the cost per viewer of running such advertisements. A very simple example of such is continuing to show bicycle-related commercials to those who haven't turned the channel when bicycles have been shown in the past, or have recently bought or searched online for a bicycle.

Thus, it can be seen that the present invention represents significant improvements over the prior art. Not only can the present invention collect more reliable data through its use of private, non-invasive data collection techniques, but the present invention can also provide data which lends itself to more advanced, thorough, and privacy-secure analysis techniques. The present invention can also analyze data more accurately than data analysis and data mining techniques of the prior art. Further, the present invention allows data analyses to be performed on behaviors observed from a larger portion of the population than the prior art, and can more accurately extrapolate such data to the population in general.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram providing a general overview of the consumer data acquisition, prediction, and query system of the present invention.

Figure 2 is a block diagram of the market data acquisition, prediction, and query system of the present invention.

Figure 3 is a block diagram of a Tuner Data Collection System of the present invention.

Figure 4 is a block diagram of a Past Events Query System of the present invention.

Figure 5 is a block diagram of a Graphic System of the present invention.

5 Figure 6 is a block diagram of the Individual Behavior Determination System of the present invention.

Figure 7 is a block diagram of the Future Events Query System of the present invention.

Figure 8 is a block diagram of the Program Entry and Program Builder systems of the present invention.

10 Figure 9 is a block diagram of the Data Mining and Prediction System of the present invention.

Figure 10 is a block diagram illustrating aspects of a sample IDM Calculation Algorithm of Figures 2, 4, 5, 6 and 9.

15 Figure 11 provides a high-level view of a technology infrastructure employed in a preferred embodiment of the present invention.

Figure 12 is a graph of linear equation values and weights for a given geographic region.

Figure 13 is an additional graph of linear equation values and weights for a given geographic region.

20 Figure 14 is an alternative view of the graph of Figure 13.

Figure 15 is an alternative view of the graph of Figure 14, and includes additional details.

Figure 16 is a sample, single-peak graph of linear equation values and weights for a given geographic region.

25 Figure 17 is a sample, two-peak graph of linear equation values and weights for a given geographic region.

Figure 18 is a sample, random pattern graph of linear equation values and weights for a given geographic region.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

Figure 1 is a block diagram providing a general overview of consumer data acquisition, prediction, and query systems of a preferred embodiment of the present invention and their interaction with each other. In this embodiment, the present invention may monitor user behavior while a user experiences television, radio, Internet, or other content. Examples of such content can include television shows, radio shows, music, advertisements, news, weather, and other multimedia or sensory-stimulating material.

Airings Data 110 comprises detailed content attributes. Examples of such content attributes include times at which such content was available; geographic or other regions to which such content was made available; actors or models appearing in or otherwise associated with such content; types of characters portrayed by such actors or models; content authors, producers, and directors; content genres, subjects, and settings; background music tones, tempo, and related characteristics; visual effect speed, colors, pixel change ratio, brightness, and other related characteristics; scents and tastes associated with such content; and other such content attributes. Additional content attributes stored by Airings Data 110 may include general plot themes or plot styles, such as comedic or dramatic segments; and a time- or position-based order in which such content attributes appear within such content.

As users view, listen to, or otherwise experience content, user behavior may be monitored through a set-top box, personal computer, radio, portable music player, or other device (“set-top box”). In a preferred embodiment, user behavior may be monitored by recording set-top box events. Tuner Data 120 may comprise a collection of such user behavior information. Tuner Data 120 may also comprise other user information, such as, but not limited to, billing information and personal demographic information.

Airings Data 110 and Tuner Data 120 may be transmitted to a data center of the present invention via a telecommunications infrastructure. In a preferred embodiment, such telecommunications infrastructure may include cable television systems, satellite

television systems, telephone systems, or other wired or wireless telecommunications systems.

In addition to the above-described information, the present invention may include data from external sources, as indicated in Figure 1 by Graphic Data 140. Graphic Data 5 140, also referred to as DATA2, may include demographic, geographic, sales, weather, and other information (“demographic information”). Geographic information used by the present invention may include, but is not limited to, distances between zip codes, population sizes within a zip code, and terrain types (coastal town, metropolitan area, etc.) within a zip code. Demographic information used by the present invention may 10 include, but is not limited to, age, race, gender, and income distributions within a zip code or sub-zip code. In a preferred embodiment, such data may be of high enough resolution as to provide information at the zip code or sub-zip code level. The only requirement is that such data share a common aspect, such as zip code, with Tuner Data 120 or any other DATA1 data with which it may be correlated.

15 Data Center 130 represents a database or other data storage device in which data from Airings Data 110, Tuner Data 120, Graphic Data 140, and the like can be stored. Data Center 130 may correlate data stored therein (such as correlating Airings Data 110 to Tuner Data 120), and such correlations may indicate content that was viewed, listened to, or otherwise experienced (“viewed”) by a user or group of users, and any reactions 20 thereto.

In a preferred embodiment, Algorithms 150 may use one or more statistical methods to determine correlations among and between data stored in Data Center 130. Such correlations may include, for example, which persons and groups viewed certain content, and why. Past Events Query System 200 may allow users to extract meaningful 25 and directed information from Data Center 130 using Algorithms 150 as described above. Past Events Query System 200 may focus on extraction of probabilities for events that have already taken place. These events can include, but are not limited to, past viewing habits of demographic groups, past sales based on advertising, past sales and viewing based on weather, and the like.

In a preferred embodiment, Past Events Query System 200 may comprise a web-based system that allows customers to query a database of past viewing behaviors. A web-based system may utilize a natural language, graphical, or command-line input interface for such queries. The Past Events Query System 200 may allow a customer to 5 query Data Center 130 while preventing a customer from obtaining any information about individual consumer behaviors, or allowing a customer to duplicate processes employed in Algorithms 150 to produce such information.

As used in this specification, the term query can include any data-related question asked of the present invention. A query may concern specific content, content portions, 10 content combinations, or mixtures thereof. By way of example, without intending to limit the present invention, a query may be, "How many African Americans in Florida, but not in Bradenton, watched at least 20% but no more than 80% of the primetime Friends last week, but did not watch at least 45% of the rerun of Seinfeld right before it or at the same time 3 weeks before that."

15 While Past Events Query System 200 can generate statistics for television airings or other events that happened in the past, the present invention is not limited to such queries. Customers, illustrated by Block 210, can also enter predictive queries through Future Events Query System 190. Future Events Query System 190 can, in turn, parse such requests into terms useable by Prediction System 180.

20 Prediction System 180 may predict future viewer behavior based upon trends found in Data Center 130 and extracted using algorithms in Algorithms 150. While these trends are of limited predictive use, such trends can be analyzed against demographic data specific to each monitored set-top box, thus providing better analysis of viewing trends across various demographic groups. Such analysis can be performed by Individual 25 Behavior Determination System 160. The addition of Future Airings Data 170 allows further predictive refinement, as Future Airings Data 170 provides a basis onto which behaviors can be mapped or extrapolated.

Customers 210 of the system may access data that has been analyzed by Algorithms 150 through Past Events Query System 200 and Future Events Query System 30 190. By these systems, a customer may tailor queries so that Algorithms 150 or

Prediction System 180 may answer them. Due to possible privacy concerns, Customers 210 may not have direct access to Individual Behavior Determination System 160, thus restricting access to behaviors of sampled households or persons.

Figure 11 provides a high-level view of a technology infrastructure employed in a 5 preferred embodiment of the present invention. As illustrated by Figure 11, analog or digital set-top boxes (Blocks 1100 to 1102) reside in viewer's homes, and can control content presentation. Such control can include, but is not limited to, the selection of a television channel or increasing or decreasing volume. Such set-top boxes may also include software which provides additional set-top box functionality, such as, but not 10 limited to, managing communications between a set-top box and a head-end (Block 1103), monitoring set-top box events, forwarding events to a head-end, and managing bandwidth utilization via configurable application parameters.

A head-end bunker can house equipment that distributes content downstream to a group of households. In a preferred embodiment, a head-end bunker can also include a 15 combination of hardware and software that monitors user behavior information from downstream set-top boxes (Block 1105). In a preferred embodiment, such set-top box data can be transmitted to a head end through a cable television cable, telephone line, or other telecommunications infrastructure. Such transmissions can also occur through a cable shared with return path equipment, even though such equipment may be separate 20 from distribution equipment.

In a preferred embodiment, a head-end bunker, as used in the prior art, may be enhanced with the addition of a UNIX-based server (Block 1106) that is connected to return path equipment via a telecommunications infrastructure. Such a server may allow collection of user behavior information.

25 A preferred embodiment of the present invention also provides a server with access to information from a customer billing database (Block 1104). Such billing system access can provide correlations between set-top boxes and customer data, such as billing zip code, billing area code and prefix, and the like. To address privacy issues regarding viewership, a preferred embodiment of the present invention will identify set- 30 top box data by zip code, area code and prefix, or other geographic identifier associated

with a region in which a set-top box resides. Correlations between set-top boxes and zip codes can be maintained in a cable television or other content provider's billing system; thus, access to such billing data may be preferred.

A highly available and highly reliable server is preferred for set-top box event monitoring, as such devices may reside in a rack at a head-end bunker, and head-end bunkers may be physically disparate or in remote regions. A preferred embodiment of Server 1106 includes a UNIX-based server; a UNIX-based server is preferred as such servers may reduce maintenance requirements. In addition, backup circuits may be implemented to provide fault tolerance depending on availability requirements for gathered data.

Server 1106 can also attach to a network access device (Block 1107) to upload data gathered from set-top boxes to a data center. Such network access devices can include, but are not limited to, modems, routers, and satellite transceivers. As illustrated in Figure 11, a private network link (Block 1108) is preferred for connecting a server to a data center for data uploads, as well as network and systems management, and for other functions. However, such functions may also be accomplished across a shared network, such as the Internet. Data transmitted across public or private networks may be encrypted or otherwise encoded to reduce the likelihood that such data may be used by unauthorized individuals.

Data uploads may occur in real-time or data may be temporarily stored on a server and transmitted to a data center on a periodic basis. Such periods may be time based, or may be based on the occurrence of an event, such as, but not limited to, receipt of a certain quantity of data or data from a particular set-top box.

In a preferred embodiment, data transmitted by Server 1106 may be received at a data center. Such a data center may be a central repository for all data gathered from a plurality of head-ends. Figure 11 includes an illustration of major data center components.

Data from Server 1106 may come into a data center through wide area circuits (Block 1109) and into temporary storage space (Block 1110). Any data cleansing or pre-processing prior to import of such data into the main database can be accomplished as

data is stored thereon. Pre-processed data may then be imported into a main data store (Block 1114).

In addition to user behavior data, a data center's main database may also store or access data from sources external to the present invention. Such external information 5 may include, but is not limited to, content attributes from various content providers (Block 1111), demographic information from third party providers (Block 1112), and sales data from retailers or producers.

As illustrated in Figure 11, data stored in a main data store may also be replicated to one or more databases for other purposes. In a preferred embodiment, data may be 10 replicated to a database that is dedicated to Internet access (Block 1113), and another database that is dedicated to report generation (Block 1115). Such replication may provide data security, as data stored in one database can be compared against data stored in other databases to ensure its authenticity. Data storage and retention properties can also be adjusted for each server as needed.

15 By way of example, without intending to limit the present invention, Internet Database 1113 can be configured to provide dedicated access to a time-limited amount of viewership data. A web-based application (Block 1117) can then provide customers with access to data in Internet Database 1113, and can also analyze and report on such data. Web servers (Blocks 1121 through 1123) can provide a front-end query system for 20 customizing such analyses and viewing reports. In a preferred embodiment, fields and data made available through Internet Database 1113 will also be structured to ensure that queries complete in a reasonable period and that impact on other users is controlled.

As an alternative example, again without intending to limit the present invention, data replicated to a reporting database (Block 1116) can be used to create hard copy 25 reports (Block 1118), electronic reports (Block 1119), and CD-ROM's (Block 1120) for customers who request access to data by means other than through a Web interface. A reporting database may also have time-limited data retention.

Figure 2 is a detailed block diagram of market data acquisition, prediction, and query systems of a preferred embodiment of the present invention. Although Figure 2 30 includes language specific to this preferred embodiment, the principles of the present

invention are also illustrated there, and can be seen with respect to any arbitrary data by replacing 'Tuner Data Center' by 'DATA1', 'Interval Updating Graphic Database' by 'DATA2', and 'Sales Data' by 'DATA3'.

While a preferred embodiment of the present invention applies the concepts of the present invention to a television ratings system, the present invention has other applications as well. Such applications include, but are not limited to, Internet advertising and actuarial analysis in the insurance industry. Whenever multiple datasets exist and correlations are desired between such datasets, the present invention can draw such correlations provided at least one dataset is relatively static and a common aspect, such as a zip code, is shared between the datasets. In the preferred embodiment disclosed in this specification, DATA1 can represent a variable data set, such as data from a set-top box, and DATA2 may represent a relatively static dataset, such as demographic data for a given geographic region. The system and method described herein can determine correlations between such datasets without direct knowledge of DATA2 values for DATA1 data.

Figures 3 through 9 illustrate individual modules of the present invention, and Figure 2 illustrates the modules of Figures 3 through 9 overlaid atop one another and interconnected, thereby illustrating interoperability of various modules and relationships between such modules. Elements of Figure 2 will be described below in connection with Figures 3 through 9.

Figure 3 is a block diagram of a Tuner Data Collection component of the present invention. In the more general functionality provided by the present invention, Figure 3 illustrates the acquisition of some data, DATA1. As it relates to a preferred embodiment, Figure 3 is a flow chart of modules useful in moving tuner data to Data Center 130 of Figure 1.

Set-Top Boxes 310 can comprise one or more set-top boxes, which can be located in one or more households. Set-Top Boxes 310 may collect and record event information based on behavior of one or more sampled users, as well as embedded content attributes, where such attributes are available. Such embedded content attributes can allow the present invention to quickly match data about specific content to set-top

box events, rather than pulling such attributes from external data sources. Embedded content attributes may pertain to content with which said attributes are transmitted, or embedded content attributes may pertain to previously presented content or content to be made available in the future.

5 Set-Top Boxes 310 may constantly transmit state-change information to HeadEnd Bunker 320, or Set-Top Boxes 310 may send batches of state-change information to HeadEnd Bunker 320. HeadEnd Bunker 320 can forward such state-change information, along with content attributes (Block 330), to Sorting System 900.

Sorting System 900 may comprise one or more sorting algorithms that place set-top box event data into efficient arrays. Due to reliability issues associated with data from set-top boxes operated by a consumer who knows he or she is being observed, these sorting algorithms may separate data into two or more classes based on whether a set-top box owner or operator has specifically requested access to monitored data, or is otherwise aware that they may be monitored.

10 In a preferred embodiment, the present invention may not collect individual-specific information, such as name, size of family, name or type of business, address, and the like from sampled users, with the exception of a zip code, area code and prefix, or other geographic identifier. Local Cable Provider 340 or other entity, which may be acting as a privacy guard for a sampled population or a governmental agency, may also 15 provide these geographic identifiers.

While data acquired by HeadEnd Bunker 320 may contain embedded content attributes, not all content may be so encoded. Non-embedded Program Information 350 may be acquired from a content airing source, such as Local Cable Provider 340, or possibly other sources, such as Internet-based guides. Non-embedded Program 20 Information 350 may comprise information that identifies content for which attributes are not available. Where such data is not electronically available, employees may make phone calls, consult published guides, or otherwise obtain such data through manual methods.

These latter methods and data collected thereby are illustrated in Figure 3 as 25 Airings Source 360. Airings Source 360 may also include a list of content that may be

available at a time in the future. In addition, Production Team 370 may work with content creators to provide content to Local Cable Provider 340 and provide Non-Embedded Program Information 350 to Airings Source 360. Production Team 370 can include an organization working with content creators who have access to program details and airings times for Non-Embedded Program Information 350. In its representation in the figures, Production Team 370 may refer to any unit or process involved in content creation or distribution, such as writers, producers, studios, networks, and the like.

As Sorting System 900 receives such data, appropriate sorting may occur and correlations may be drawn between such data and data from Block 330. Sorted data may be stored in Tuner Data Center 930. Tuner Data Center 930 may comprise a database of set-top box data arrays of relevant age. Arrays of information of non-relevant age may be stored off-line, but may still be permanently accessible.

Figure 4 is a block diagram of Past Events Query System 200 of Figure 1. Although illustrative of a preferred embodiment of the present invention, Figure 4 also illustrates the general concepts of the present invention with respect to any arbitrary DATA1 and DATA2 if ‘Tuner Data Center 930’ is replaced by ‘DATA1’ and ‘Interval Updating Graphic Database 620’ is replaced by ‘DATA2.’

Graphic Vendor 610 may comprise one or more data vendors supplying the present invention with demographic data for a geographic or other region. Graphic Vendor 610 may provide such information broken into distribution units, where such distribution units share a common factor such as zip code or sub-zip code.

While Graphic Vendor 610 may supply such data in a preferred embodiment, an alternative embodiment can replace data supplied by Graphic Vendor 610 with data determined internally by a Graphic System as illustrated by Figure 5. A Graphic System can use data collected by an Individual Behavior System, which is illustrated in Figure 6, to provide necessary data. In another embodiment, data from a Graphic System can be augmented by data from Graphic Vendor 610 to provide the present invention with more comprehensive data.

In a preferred embodiment, the present invention may periodically request data from Graphic Vendor 610 and such data can be stored in Interval-Updating Graphic Database 620. As Interval-Updating Graphic Database 620 receives data from Graphic Vendor 610, data stored in Interval-Updating Graphic Database 620 may be modified to reflect changes implied by data from Graphic Vendor 610. Interval-Updating Graphic Database 620 may then be used by the present invention as a source of graphic data.

The present invention may use data from Interval-Updating Graphic Database 620 to create data arrays representing relative data distribution percentages. Such arrays can be compiled by IDGM Graphic Matrix 910 at any time. In a preferred embodiment, 10 IDGM Graphic Matrix 910 may be updated when new data is received by Interval-Updating Graphic Database 620.

IDGM Graphic Matrix 910 may create matrices for each set of graphic data. Such matrices may contain arrays that refer to zip codes or other geographic descriptors to which information contained within an array corresponds. In a preferred embodiment, 15 arrays may be formed with column headings corresponding to graphic characteristics, such as, but not limited to, gender or age, and rows corresponding to a set of zip codes. A number corresponding to the percentage of said row that can be attributed to said column may be stored in the intersection of each row and column. Thus, for example, if 65 percent of the population of a particular zip code were male, 0.65 could be stored in 20 the intersection of the male column and the row corresponding to said zip code. Such arrays can then be used by a Process Computer for matrix operations that provide numerical data to a report processor prior to delivery to a customer.

Customers, illustrated in Figure 4 by Market Customer 530, may request such reports from Past Events Query System 200. The present invention may translate such a 25 request into a mathematical formula, or a machine-language representation of such, through Post-Translation System 280. Formulae created by Post-Translation System 280 may be interpreted by IDGM Calculation Algorithm 270 to properly extract and analyze data stored in IDGM Graphic Matrix 910 and Interval-Updating Graphic Database 620.

Figure 10 provides an overview of a sample IDGM Calculation Algorithm 270 30 that can perform such analyses. As illustrated by Block 1012, algorithms used by IDGM

Calculation Algorithm 270 may take advantage of an assumption used by media researchers, which is that each member of a given DATA2 (Block 1011) group has the same probability of exhibiting some DATA1 (Block 1010) behavior as any other member of the same group. The present invention extrapolates from this an assumption that
5 probabilities associated with behaviors of groups of people can be determined by their demographic specification (this is referred to as the "demographic assumption"). IDGM Calculation Algorithm 270 uses calculations derived from these assumptions to develop DATA2 correlations for DATA1 data without collecting DATA2 information about DATA1 data directly, and to determine confidence intervals associated with such
10 correlations.

IDGM Calculation Algorithm 270 may use inverse mathematical principles to find such correlations (Blocks 1013 and 1014). The following are two methods through which such correlations may be determined by IDGM Calculation Algorithm 270. While these examples are provided for enablement and best mode purposes, these examples
15 should not be construed as limiting the present invention. In alternative embodiments, subsets of these examples may be used, as may additional calculation methods.

A preferred embodiment of the present invention assumes that a person's demographic description has some influence on his choice of television viewing. A goal of this embodiment of IDGM Calculation Algorithm 270 is to apply the inverse of this
20 assumption; that is, a person's demographics can be determined from his viewing habits. To achieve this, the present invention can invert region-specific viewing and demographic data to compute demographic-specific viewing information.

The following are definitions that will aid in understanding this embodiment:

Demographic Data – As with this specification as a whole, the term "demographic data" includes ordinary demographic categories as well as geographic variables and general local characteristics. Examples of ordinary demographic categories include, but are not limited to, age, race, gender, income, education level, marital status, and number of dependents. Examples of geographic variables include, but are not limited to, climate and weather, urban or rural environment, coastal or inland geography, population density,
25

and amount of traffic. General local characteristics may include, but are not limited to, progress of regional sports teams and local news events.

Demographic Characterization - A demographic characterization is a set of values for each of a given set of demographic categories.

5 Demographic Characterization Level – A demographic characterization level is the number of categories comprising a demographic categorization. For example, a level-one characterization might be a specification of race, while a level-two characterization might be an age group together with a race.

10 Demographic Specification - A demographic specification is a full demographic characterization that uses all demographic categories tracked by the present invention.

Demographic Aspect - Demographic aspects are potential demographic category values. For example, the category “gender” has aspects male and female.

15 Orthogonal Characterizations - A set of demographic characterizations is said to be orthogonal if there is no overlap among them; that is, a given person can fit into no more than one of them.

Complete Characterizations - A set of characterizations is said to be complete if any given individual person necessarily falls in at least one of the characterizations in the set.

Orthogonal and Complete Characterizations - A set of characterizations is said to be orthogonal and complete if any given individual falls into exactly one characterization.

STB - A set-top box.

25 Program State - Program states reflect particular content, or portions and combinations of content, presented by an STB at a particular time. For example, a program state could be defined as some specific 10-second interval of a particular commercial which just aired, combined with an entire program that aired 2 weeks ago.

Tuner State - Tuner states represent a current STB state.

Event Rating – An event rating represents a probability that a person or group of people, represented by a demographic characterization, matched or will match an STB event.

ERINRating – An erinRating is produced for each event rating by prorating all 5 event ratings for a given geographic area and over a given time period with a given set of program state choices. The number of persons in a geographic area exhibiting an event can be determined by multiplying the number of persons matching a demographic characterization in a given area by an associated erinRating.

STB Event Time – STB event time is a time series that is defined by STB event 10 sampling. For example, if a cable provider's system is inoperable for 10 hours, this gap is not considered in STB event time.

A goal of the present invention is to determine the extent to which the demographic assumption is valid. This determination can be a factor in calculating confidence intervals for data resulting from the present invention.

15 The demographic assumption can be expressed mathematically in the following relation equation:

$$m_k = \sum_i p_{ki} v_i$$

Equation 1

Here m_k is an observed number of STB's in zip code k experiencing some defined 20 event, p_{ki} is the number of people in zip code k with demographic characterization i , v_i is the fraction of people of characterization i that are watching the event, and the sum is over a complete set of demographic characterizations i . This formula embodies the demographic assumption because v_i depends only on i . In English, Equation 1 simply says that the total number of STB's experiencing a particular event can be determined by 25 summing the number of people of each demographic characterization that are experiencing the event.

In a more general application of the present invention and its related formulas, m_k can be seen as corresponding to DATA1, and p_{ki} to DATA2.

Given that m_k values can be determined through the present invention, and p_k values can be obtained from demographics vendors or by other means, the present invention can use Equation 1 to solve for v_i . This can be accomplished by defining an error functions such as:

5

$$\psi^2 = \sum_k \sum_i (m_k - p_{ki} v_i)^2$$

Equation 2

If a dataset under consideration contains more zip codes than demographic characterizations, Equation 2 can be solved through a standard least-squares approach. If a dataset under consideration contains fewer zip codes than demographic

10 characterizations, fitting methods may be applied prior to application of a standard least-squares approach. A least-squares approach can involve inverting a matrix based on p_{ki} , and for that reason is referred to as an inverse demographic matrix, or IDM, solution.

This is illustrated by Block 1014 in Figure 10.

An IDM can be implemented in different manners, depending on the set of 15 demographic characterizations. For example, if a particular query involves only one category, such as age, then demographic characterizations can be defined across the whole set of age-intervals (0-10, 11-20, 21-30, ..., 101-110, 111+, for example). In this case, a set of twelve age-intervals forms a complete (all individuals fall into at least one interval), orthogonal (an individual falls into no more than one interval) set of 20 characterizations. A resulting p_{ki} matrix will then be a matrix of size $N_{zip} \times 12$, where N_{zip} is the number of zip codes used in the calculation. Equation 2 can then be solved for twelve values of v_i , v_1 through v_{12} .

Alternatively, if a query involves only one particular age-interval, such as ages 11 to 20, a set of demographic characterizations with only two elements can be used, one set 25 containing people between 11 and 20, and one containing all others. The resulting set is also complete and orthogonal, and can be represented in a p_{ki} matrix. In this case, the p_{ki} matrix is $N_{zip} \times 2$, and Equation 2 can be solved for two values of v_i , v_1 (the number of people in the age range 11 to 20 who are watching), and v_2 (the number of people of all other ages who are watching).

An IDM solution for a query involving a set of level-*n* characterizations may be referred to as IDM*n*. The age-group example given above, using an $N_{zip} \times 12 p_{ki}$ matrix, can thus be seen as an example of IDM1. If a query involves two categories, for example age group and gender, then a complete, orthogonal, level-2 demographic characterization set would have 24 elements, and an IDM2 solution would involve a p_{ki} matrix of size $N_{zip} \times 24$.

When a query involves only a single characterization, such as ages between 11 and 20, this may be referred to as IDM*n*-P. Thus, the second example in the previous paragraph is an IDM1-P solution. As another example, if a query only involved women in age interval 21 to 30, then an IDM2-P solution, for which the p_{ki} matrix would again be of size $N_{zip} \times 2$, could be used.

IDM solutions can be normalized with complementary IDM*n*-P solutions to supplementary IDM(1,2,3...n)-P solutions. Complementary normalization involves IDM*n*-P computation of all characterizations of the same “n” which contribute to a whole IDM(*n*-1) mutually exclusive demographic characterization. IDM*n*-P values can then be normalized to total the value computed for the IDM(*n*-1)-P characterization (which was previously normalized itself if *n*>1). All normalization can begin at characterizations of lowest n value. By way of example, when n equals 7, characterizations 1 through 6 should be calculated, so that normalization occurs between all two-level combinations.

While the present invention may calculate an IDM solution, such a solution may not be presented to a customer querying the present invention. Rather, the present invention may present a range of values that fall within a particular level of statistical confidence. As STB usage expands to include the majority of the population, such value ranges may have a non-zero size due to a violation of the demographic assumption. Values for such ranges can be determined by uncertainty estimations provided by a least-squares fit.

While STB's are well below complete penetration but data is sampled in a random fashion, value ranges returned as a result of a query can have an additional uncertainty contributed by sampling error. In addition, sample bias, which can occur, for

example, when a sampled individual knows of such sampling, or simply as a result of a difference between those willing to be sampled and those unwilling to be sampled, can cause additional sampling error. The methods outlined below can address these complicating issues, and can calculate error ranges for various datasets.

5 As data is collected by the present invention, a 2-dimensional array holding the number of matching events between all combinations of demographic characterizations is kept in a specification similarity matrix, illustrated by Block 1026 in Figure 10. For each demographic characterization combination, the Pearson-r correlation can be computed with respect to all pre-defined events.

10 Over time, IDM values for demographic characterizations or combinations of demographic characterizations can become predictors for demographic characterizations, within determinable statistical confidence intervals. Such demographic characterizations can be combined to create a similarity index. The present invention can use a similarity index to determine probability ranges for various levels of demographic characterizations 15 that can be used by a possibilities reduction system of the present invention.

The present invention may also apply another assumption to determine a combination of demographic aspects that can have a relationship on television viewing. Such relationships can be determined by applying rules to demographic aspects over time. Such rules may include, but are not limited to, additive, subtractive, and dominance 20 or recessiveness rules. If an IDM2 value for a level-two characterization is greater than the IDM1 value for both demographic aspects comprising the level-two demographic characterization when viewed alone, that value is said to correspond to an additive rule. If an IDM2 value for a level-two demographic characterization is less than the IDM1 value for both level-one demographic characterizations when viewed alone, that value is 25 said to correspond to a subtractive rule. If an IDM2 value for a level-two demographic characterization falls between IDM1 values for the respective demographic aspects of the level-two characterization, that behavior is deemed to be dominant/recessive, and the IDM1 value closest to the IDM2 value is deemed the dominant IDM1 value.

A multi-dimensional array can be kept which records rules appropriate to each 30 IDM relationship for each event over time. Such an array may also be extended to

include rules comparing multi-aspect combinations. Statistical tests can be applied to determine confidences with which an aspect or combination of aspects is related to another aspect or combination of aspects by any of the given rules. Weights may be assigned to each rule, with a preferred embodiment using linearly higher weights to 5 represent exponential rule growth.

Values produced by such a system may be stored in an array, or recombination matrix. Rules can be applied additively for a demographic characterization of a particular level from lower level IDM solutions at which such rules are identified. Final confidences can be determined through Pearson-*r* correlation with IDM calculations.

10 Demographic characterization recombination matrices can aid in the calculation of probability ranges.

A mean of each demographic specification's event matching representation may be stored over time in aspect representation indices. Such one-dimensional arrays can also be used in confidence determinations for information obtained through other 15 portions of the present invention. These indices can be updated as IDM calculations continue through STB event time. Aspect representations can give the present invention approximate sample sizes for behaviors of each demographic characterization.

Aspect representation indices can also be used when determining individual behaviors. To determine such behaviors, probabilities can be assigned to each 20 demographic specification and each STB event. If an STB matches an event, corresponding probabilities can be ascribed to an STB, where such probabilities are normalized so the highest probability is unity. If an STB does not match an event, probabilities ascribed to such an STB may be the linear inverse of probabilities ascribed to an STB matching an event. For each STB and demographic specification at a given 25 time, summing ascribed probabilities for each demographic specification and dividing by the number of probabilities can compute the probability that an STB corresponds to a given demographic specification. A one-way analysis of variance can then be performed on such data to determine the likelihood of such data representing a user of a respective STB.

As previously discussed, confidences with which a demographic specification can be linked to an STB can be useful in refining data generated by the present invention. This usefulness arises out of the fact that combining relative confidences with a level-n demographic characterization yields a level-n+1 demographic characterization. Such 5 relative confidences are also useful when evaluating assumptions, or rules, generated by the present invention. Such assumptions may be generated in the above-described aspect recombination rules, specification similarities, aspect representation, and individual behavior determination processes.

Information developed by each assumption has an empirical validity, which can 10 be converted to a statistical confidence. This empirical validity can be determined over STB event time by the assignment of expectation values, tracking of empirical values, and through time-correlation tests between expectation and empirical values. Each assumption's validity may require determination through a slightly different statistical formula.

15 Assumption validities, along with sample numbers and sample values, can provide probabilities and ranges for information concerning demographically specific groups. The present invention may translate information into a chosen confidence interval, then, for each specification and each process performed, sets of individuals matching and not matching all possible events can be produced.

20 Through this process, each demographic specification group can be labeled with statistical "guesses" at a final range of values for an event rating. A system of linear equations can then be solved to further reduce the ranges and fill in residual gaps left by such processes. Demographic numbers for each demographic characterization can also be used to further reduce the ranges by this possibilities reduction system. By way of 25 example, without intending to limit the present invention, imagine a set of viewers matching a set-top box event contains 300 people who are Asian, and this same set contains 200 people who earn over \$80,000 per year. If it is known that all people in the set who earn over \$80,000 per year are Asian, then the 100 remaining Asian people can be placed into categories corresponding to incomes below \$80,000 per year. Some of 30 these categories may already be reduced, so information can be filled in quickly.

Essentially, while an IDM and other processes of the present invention may have increasingly low confidences as demographic specificity levels increase due to low sample representation for many of demographic characterizations at these higher levels, a possibilities reduction process can take advantage of these low numbers by filling in 5 demographic characterizations with certain matches or non-matches of STB events. A possibilities reduction system can accomplish this because demographic data is known at these levels, and thus reduces the set of possibilities remaining.

Ranges can be reduced even in cases where demographic characterization numbers are not filled completely. The possibilities reduction system may use linear 10 algebra rules to describe the remaining possibilities in terms of mathematical symbols after each piece of information is considered. A matrix of specific probability ranges and their mathematical relationships with other specific probability ranges for a given event can be iteratively updated in a pre-determined order until such iterations do not significantly change any matrix values.

15 Ultimately, ranges that cannot be further reduced may not be modified by this iterative process without first modifying the confidence level. For example, by decreasing the confidence level (for example, from 90% to 88%), ranges may be reduced and some demographic characterizations be filled in. An iterative process can then go through all characterizations and iteratively again reduce value ranges as much as possible. At some 20 point, the confidence level in the matrix may be increased to its original level. However, additional mathematical procedures may be introduced to avoid an error function of the of this matrix getting trapped in a (spurious) local minimum.

It is important to note that even if a customer is only satisfied with confidences 25 associated with level-five data, the relevant level-five characterizations can be composed of any five demographic categories about which the present invention collects demographic data and in which a customer may be interested.

In addition to a probability reduction process, a preferred embodiment of the present invention may further reduce the likelihood of errors by employing a Monte Carlo-type fit to such data after a probability reduction process has been applied. Such a 30 fit may account for bias issues associated with data collected by the present invention.

The present invention may further address bias issues by monitoring set-top box usage and identifying inconsistent behavior. For example, the present invention may detect when households leave a set-top box on while playing a DVD, videotape, or even while on the telephone or when out of the home for extended periods by analyzing usage patterns for each set-top box. By way of example, without intending to limit the present invention, the present invention may define “special events” which correspond to situations in which a particular set-top box remains on but no state changes occur over a period that is significantly longer than an average state-change interval for that set-top box.

10 The present invention may also be integrated with recording systems, such as the Replay-TV and Tivo systems, to allow more detailed analysis of consumer behavior. Recording systems may be of interest in the present invention because of the level of control a user has over a given program, including the ability to pause live programming.

15 An alternative calculation method which may be used by IDGM Calculation Algorithm 940 essentially involves a more additive, yet thorough approach based on the demographic assumption to determine DATA2 values for DATA1 data. However, rather than defining an error function to be minimized, this method represents a lengthy process of linear algebra by essentially comparing each zip code to each other zip code for each demographic characterization in a tailored customer query. This calculation method is an
20 alternative mathematical method that solves the same problem as the previous method.

1 The present invention may essentially translate customer queries into a query of the type “What percentage of a certain demographic specification performed some action during a program state?” and calculate the result of such queries. Definable Viewing Units (“DVU’s”) include whole or partial content that can be matched to time and
25 location. Customer queries can take the form of combination of actions and DVU’s.

Customer queries can be translated to a final query through determination of a plurality of values within subsystems of the present invention and then operating on those values to derive a final query value. Subsystem query values can be determined by a process computer, which employs multiple sub-processes to determine each sub-value. A
30 processor can convert any percentages generated by such sub-processes to numbers,

prorate and extrapolate such numbers to include non-sampled individuals, and accounts for those STB's at presenting content to multiple viewers.

A process computer may receive queries specific to a demographic category. A process computer can break apart a query into two components; a set of zip codes for which the query will return data, and all other query components. For each zip code in a query, which may be referred to as a "query zip," the query zip can be matched against other zip codes for which the present invention collects market data. This process will result in n-1 zip code combinations ("zip-zip combination"), where a query zip is the first combination, and where n is the total number of zip codes about which data is collected by the present invention. These zip-zip combinations can be run through a multi-step process that determines final query values.

The first of these processes, CP1, determines weights for each zip-zip combination. CP2 determines a value each zip-zip combination can infer for the final query, or for some sub-query. CP3 can determine if there is more than one pattern of inferred query values and, if there are, determines a set of values from the current dataset. CP4 reduces values from CP3 by one level. CP3 and CP4 can be iteratively acted upon until only one pattern is apparent, and a result set can be determined.

CP1 determines weights for each zip-zip combination based on a variety of factors. One such factor is the percentage of a test zip code's population that falls into the query category in question. In a preferred embodiment, a test zip code's population should be as different as possible from a query zip. This provides a resolution increase, as the larger the differences between the categories, the better the effect of individual common points can be determined.

For each category other than zip codes, CP1 may give higher weight to those categories for which the population of the test zip is similar to the query zip. The importance of these other categories can be determined to make an overall determination of how similar the test zip is to the query zip in terms of such other categories. By way of example, without intending to limit the present invention, if test zip(1) is ninety percent similar by age distribution and ten percent similar by religious affiliation, test zip(1) will likely be assigned a higher weight than one that is ten percent similar by age

distribution and ninety percent similar by religious affiliation. The relevance of each category, and hence its weight, can be determined by the effect a combination of zip codes seem to have on market share for a given query.

The present invention can determine a weight for a given test zip by the formula
5 Weight = $(c)(W)$, where c is the percent difference between persons belonging to a query category in a test zip and those belonging to the same query category in a query zip.

As an example, without intending to limit the present invention, if CP1 received a query including query category 'African American', and the following were true for sample query zips and test zips:

10 query zip: 10% AA + 90% other = Market Share = 20%

test zip: 20% AA + 80% other = Market Share = 40%,

then c may be calculated as a percent difference between 10.0 and 20.0.

W can be determined by summing weights assigned to each "other" category, where other categories are defined as categories in the union set of categories between a 15 query zip and test zip, except for a query category or any sub-categories defining a query category. In equation form, this can be written as $W = \text{sum}(w[x]) = \text{sum} (w[1] \dots w[x])$, where x is a numerical identifier of each category, and x is incremented over the number of categories.

For each category:

20 Let $\%d(x,y) =$ the absolute value of a percent difference between 'x' and 'y'

Let $a = \%d(A[0], A[n])$, where $A[0]$ represents a demographic percentage in query zip for a given demographic category, and $A[n]$ represents a demographic percentage in a test zip for a given demographic category.

25 Let $b = \%d(B[0], B[n])$, where $B[0]$ represents a market share for a given demographic category in a query zip, and $B[n]$ represents a market share for a given demographic category in a test zip.

Let each category weight $w(x) = q(a) \times q(b)$

where $q(a) =$ the weight of 'a' and $q(b) =$ the weight of 'b' where:

$q(a) = f(a) = \text{approximately } 1/a$

$q(b) = f(a,b) = \text{a function in which:}$

as 'a' goes to 0 and $\%d(a,b)$ goes to 0, $q(b)$ goes to infinity, and

as 'a' goes to infinity and $\%d(a,b)$ goes to 0, $q(b)$ goes to 0.

5 For example, if a query category is African American, categories in a complete union set of categories between the query zip and test zip which do not include African Americans should be reviewed. One such category may be White Male, and by way of an example, the following may be true:

query zip: 24% WM + 76% other = Market Share = 20%

10 test zip: 30% WM + 70% other = Market Share = 40%

The statistics above can be rewritten as follows if the "other" categories and their respective percentages are disregarded:

query zip: $A(0)$ WM = $B(0)$, where $A(0) = 24\%$, $B(0) = 20\%$

test zip: $A(n)$ WM = $B(n)$, where $A(n) = 30\%$, $B(n) = 40\%$

15 An optimum weight function for a test zip code is an intrinsic property of its relationship to a query zip. $q(a)$ can be seen as a measure of a percent similarity of a given category between query and test zips, and $q(b)$ can be seen as a measure a category has on describing market share differences between query zip and test zip. Thus,
20 categories receiving heavier weights are those in which $\%d(A[0], B[n])$ is low or null and at these low values $\%d(B[0], B[n])$ is a similar trended value. In a preferred embodiment, As $\%d(A[0], B[n])$ increases, that $\%d(B[0], B[n])$ should decrease.

25 A natural function fitting $q(a)$ and $q(b)$ requirements will be one that provides a proper optimization of zip codes, such as a Laplacian function. It is clearly a symmetrical surface in 3D space with four specific boundary conditions. If a specific function can not be found that provides such a fit, a practically optimal function can be created through power series.

CP2 can evaluate the percentage of persons falling into a query category in a query zip and the percentage of persons falling into the query category in all test zips, and

use this to evaluate specific market share differences between the query zip and each test zip codes. CP2 then evaluates the information about these two zip codes, and establishes a best guess as to the query category differences that contributed to an observed market share difference.

5 CP2 can take demographic and market share information corresponding to individual zip codes and solve a set of linear equations for variables involved. One variable resulting from such a solution can be related to a percentage of persons falling into a query category for a given zip code or set of zip codes. Another such variable may be the percentage of persons in a zip code not falling into a category. CP2 may set the
10 right-hand side of each equation to the market share of a query DVU.

By way of example, without intending to limit the present invention, a given zip code or set of zip codes may yield a formula such as: $10AA + .90o = .20$. Such a formula may indicate that a zip code or set of zip codes has a demographic makeup that is ten percent African American and ninety percent "other." To further ease understanding of
15 this example, assume the DVU is the percentage of TV sets that watched all of the show ER from 10.00.00 to 11.00.00 on 11/11/99. In the above example, the market share for this DVU is 20 percent. It should be clear to one skilled in the art that in the equation above, the variable 'AA' refers to the percentage of African American persons in a given zip code fulfilling DVU criteria.

20 A preferred embodiment of present invention may represent such an equation using the following three matrices for computational clarity and storage efficiency:

$$[AA \ o] \quad [10 \ 90] \quad [20]$$

CP2 may translate information for all query and test zip codes into this form for a query category and for other categories in which none of the persons making up those
25 category percentages fall into a query category. If the zip code, category, and DVU example above were expanded to include a test zip(n), there would be 3 matrices again, but this time 2-dimensional:

$$[AA \ o] \quad [10 \ 90] \quad [20]$$

$$[AA \ o] \quad [12 \ 88] \quad [19]$$

Relationships between such matrices and their resulting linear equations can be seen when the linear equations are represented graphically by a line in 2-dimensional space. If two lines are on the same plane, such lines must either 1) be the same line, 2) be parallel lines, or 3) cross at some point. Through these relationships, CP2 can create a 5 "best guess" at a mutual answer for variables in the equations based on information the equations imply. This guess need not fall onto a point in the interval of either equation.

This process can be seen as analogous to drawing a set of lines in 2-dimensional space and then finding intersections of these lines. This intersection may occur once per 10 two dimensions, indicating two particular variables. The first of these variables relates to a particular category of interest, and the second variable is always 'other'.

Statistically, with several test zips with which to work and each with a given number of samples, CP2 may be seen as analogous to sampling a much larger number of samples than are in any individual sample zip, and drawing a normal curve for each. The 15 resulting curves may then have each curve subtracted from them, and one normal curve may be built from the differences. In some cases, this curve may be a line representing a single value, rather than a curve.

At this point, the normality of an extrapolated curve is no longer an assumption, as only one specific category variable is of concern at a given time. Any "other" values 20 not under the peak of a curve represent a mathematical confidence level based only on sample size versus population error, and not a statistical skewing possibility.

CP3 can take information from CP1 and CP2 in the form of a series of test zip codes and the best guess value for each. Based on this data, CP3 can then determine whether such best guesses exhibit a one-dimensional pattern. If such a pattern exists, CP3 can return the peak of the pattern as the query return value. If such a pattern is not 25 exhibited, CP3 may pass data to CP4.

CP4 takes CP3 data and attempts to establish a one-dimensional pattern from it. CP4 can run multiple queries through CP1, CP2, and CP3 using each zip code in the present invention's sample population as a query zip. For each iteration through CP3, a photographic category percentage can be assigned to each zip code. These can be

reinserted to CP1, CP2, and CP3 using a query zip and its assigned phantographic categories in place of demographic category percentages.

With each iteration, the complexity of CP4 data is reduced by one level. When CP4 data exhibits a one-dimensional pattern, this result can be returned as a final query value. If CP4 data exhibits more than a one-dimensional pattern, the data can be fed back through CP1, CP2, and CP3 for additional analysis, using new phantographic category percentages with each iteration. Through this process, CP3 and CP4 can identify those demographic categories having a high effect on viewership.

As previously described, CP3 can receive data from CP1. Such data may be in
10 the following form:

test zip 1	zip 1 weight
test zip 2	zip 2 weight
test zip 3	zip 3 weight
...	
15 test zip n	<u>zip n weight</u>
	total = 1.00

In addition to data from CP1, CP3 may also receive data from CP2. Such CP2 data may resemble the following:

20 test zip 1	query zip - zip 1 linear equation solution value
test zip 2	query zip - zip 2 linear equation solution value
test zip 3	query zip - zip 3 linear equation solution value
...	
25 test zip n	query zip - zip 3 linear equation solution value

CP3 can multiply solution values for each zip code by the weight of that zip code to get a sub-answer. These sub-answers can then be summed to yield a final answer.

While CP3 can generate query results based on linear equation solutions, CP3 can also generate and analyze graphs of zip code weights (w) versus linear equations solution values (v). An example of such a graph is illustrated by Figure 12.

CP3 may reduce the influence of random noise on such a graph by assigning each w the average of its value and the values on either side of it. This procedure may be repeated until all noise is reduced to acceptable levels. The result would be a graph with one of the following conditions:

- 1) one defined peak from one direction;
- 2) one defined peak from two directions;
- 3) more than one defined peak, possibly with varying heights; or
- 4) a random curve.

5 If a curve exhibits conditions outlined in numbers 1 or 2 above, a near-perfect guess at an ultimate query return value can be made. Such a value may be based on a peak value observed, or may result from extrapolation of a graph over a larger data range.

If a curve exists exhibiting conditions outlined in number 3 above, which is illustrated by Figure 12, then at least one additional category may be effecting query data.

10 A random category can be created, called a phantographic, and relevant percentages can be assigned to that category to account for peak distributions observed in the graph. CP4 can then feed this phantographic category back into CP1 and CP2 for the query in question to determine an appropriate query return value.

By way of example, if the data in Figure 12 represented a 60/40 distribution, the
15 60 percent distribution may be assigned to zip-zip combinations where test zips are directly under parts of the graph associated with the 60 percent peak. The 40 percent distribution may be assigned to zip-zip combinations where test zips are directly under portions of the graph associated with the 40 percent peak. Test zips falling under valleys between such peaks may be ignored. Through this system, the present invention should
20 account for categories affecting market share differences without requiring tracking of a large number of categories.

As CP3 feeds data to CP4, CP4 can assign phantographic category percentages to each zip code, then route the result through CP1, CP2, and CP3, thereby reducing underlying patterns by one level. The following is an example of calculations and
25 procedures employed by CP4.

Figure 13 is a sample graph generated by CP3 for a query category, zip code, and market share. CP2 values are arranged in increasing order from left to right, with higher lines indicating heavier weighting for a given value. When CP2 generates two or more of

the same value, the weight represented on such a graph may be increased to the sum of the individual weights for all matching values.

Figure 14 provides a more detailed view of the graph in Figure 13. The sum of all weight values in the graph should total 1, and all weights must fall between 0 and 1. This
5 is due to the prorating step performed after CP1. Figure 15 provides an additional view of Figure 14, with values provided for various points on the graph, and zip codes numbered 1 through 25. The following is a table of values illustrated by Figure 15:

	Zip Code	weight (to 3 significant digits)	value (to 2 significant digits)
10	1	0.402	0.50
	2	0.515	0.79
	3	0.027	0.23
	4	0.040	0.71
	5	0.038	0.90
	6	0.070	0.29
15	7	0.054	0.41

	25	0.052	0.26

In a more abstract sense, data passed from CP3 to CP4, when graphed, may
20 resemble Figure 16, Figure 17, or Figure 18. Figure 16 illustrates CP3 data resulting in a graph with a single peak. Figure 17 illustrates CP3 data resulting in two peaks. Figure 18 illustrates CP3 data with many peaks. While CP3 data may resemble one of these three figures, CP4 does not distinguish between such cases.

If, for a given query, CP3 resulted in a graph similar to Figure 16, this may result
25 in a great deal of confidence in a query result value. Such a value may be either the sum of the means of each value times its respective weight or simply the value underneath the peak of the graph. Assuming a symmetrical and "natural" weighting function or algorithm, CP3 data yielding Figure 16 would provide validity to the assumption that persons with more similar graphic makeup have more similar television viewing patterns.
30 The present invention can answer the questions customers truly intend to ask, as any quality shared by that group alone is represented in the CP1,2,3 algorithm. This alone can represent a significant improvement over statistical methods employed in the prior art.

Rather than the single-peak graph of Figure 16, if CP3 were to result in a graph
35 similar to Figure 17, the correct answer is not one that would be given by a full average

of all data, as this would provide a value between the two peaks. The correct answer is the value on the V axis directly under the first peak, or the value on the V axis directly under the second peak. CP4 can select an appropriate value from two or more such values.

5 The fact that there are two peaks in Figure 17 indicates that another demographic category should have been part of the graphic data. When two peaks exist such as in Figure 17, there is only one difference between the set of values forming the peak on the left versus the set of values forming the peak on the right. While CP4 may not determine what this difference relates to, CP3 data is based on zip codes, and the difference is likely
10 to be geographic. For example, it may be that the peak on the right is formed by zip codes in coastal cities, and the peak on the left is formed by non-coastal cities.
Regardless of the source of the difference, Figure 17 clearly illustrates its existence.

If all graphic categories are ignored except the query category, a new set of categories, called phantographic categories, can be assigned to test zips making up the
15 chart distribution. A value of 0 can be assigned to test zips creating the left-hand peak, and a value of 1 can be assigned to those test zips creating the right-hand peak. If this data set is now run through CP1, CP2, and CP3, a similar but more disperse graph should result. If all zip codes sampled by the present invention were iterated through as query zips, and all others were test zips, and the value belonging to each peak was assigned to
20 those test zips beneath it as the phantographic category percentage for each query zip in the system, this data could be run through CP1, CP2, and CP3.

This would result in a new graph for the query zip. Such a graph should result in all phantographic categories having similar percentage distributions for a given zip code. Further, a percentage exhibited by a query zip can be determined, and thus the initial
25 peak associated with the query zip can be properly selected by CP4.

The present invention may further employ a "category generator", which can search demographic, geographic, and other databases for a graphic distribution percentage matching any recurring phantographic distributions. While many are likely to remain undiscovered, if one is found it may be added to a list of categories monitored by
30 the present invention.

Due to CP4, patterns in CP3 data cannot escape the present invention. Even algorithms at the heart of data generated by a random number generator could be determined through CP4's iterative processes.

Figure 5 is a block diagram of modules used in Graphic Data acquisition. As Figure 5 illustrates, the present invention may acquire data for Interval-updating Graphic Database 620, and ultimately for IDGM Graphic Matrix910, from outside sources such as Graphic Vendor 610. However, as also illustrated in Figure 5, the present invention may also garner data for Interval-updating Graphic Database 620 from Evolving STB Viewer Possibilities Database 960.

Evolving STB Viewer Possibilities Database 960 may include, for each set-top box, behaviors and demographics associated with one or more regular users of said set-top box. Evolving STB Viewer Possibilities Database 960 may also identify a set of set-top boxes whose monitored behavior best fits demographics for a given geographic region. In a preferred embodiment, regional demographics may be complied from specification percentages of such set top boxes. In an alternative embodiment, specification percentages may be fit by the weight of their magnitudes. The present invention may correlate and assign demographics to set-top boxes, and the sum of these can be combined to indicate demographics for a specific region.

While Interval-updating Graphic Database 620 may initially receive data from Graphic Vendor 610, Interval-updating Graphic Database 620 may not always require such data. As the quality of data stored in Evolving STB Viewer Possibilities Database 960 increases, the present invention may no longer require data from Graphic Vendor 610.

Block 460 identifies a process that produces data which may raise privacy concerns. As illustrated in Figure 5, the present invention can protect such data by restricting access to such data to only components of the present invention. Information stored inside Block 460 can only be accessed by customers through components of the present invention, such as Block 940.

Figure 6 is a block diagram illustrating modules used by an Individual Behavior Determination System of the present invention to acquire individual behavior data. As

illustrated by Figure 6, IDGM Calculation Algorithm 940 may extract event data from Tuner Data Center 930. IDGM Calculation Algorithm 940 may also perform event analysis.

As with other parts of the present invention, an individual behavior system may
5 use the demographic assumption at its core. An individual behavior system may also assume that statistical determinations can be made that define a viewer of an STB by measuring STB states over time and ascribing probabilities to an STB. Individual viewers of an STB can then be statistically identified based on behaviors exhibited on an STB.

10 The mean of all ascribed specification percentages may be kept in Evolving STB Specification Percentage Database 950, illustrated in Figure 6. Evolving STB Specification Percentages Database 950 may comprise, for each specification, a percentage of time that a behavior associated with a set-top box matches a particular specification. Viewed over time, Evolving STB Specification Percentage Database 950
15 may contain a “fuzzy” percentage, or Identity Percentage, for each specification, which may be determined for each set-top box.

The “Attribution Percentage” ascribed to an STB for a given query behavior ‘n’ can be calculated through a piecewise function:

$$AP_{(n)} =$$

20 If YES: behavior match % of that specification;

$$\text{If NO: } \left(\frac{1}{\text{Process 1 \% of spec}} \right) / \sum_{1}^{s(c)} \left(\frac{1}{\text{Process 1 \% of spec}} \right)$$

Where s(c) is the total number of mutually exclusive specifications of the query
25 category / categories.

The “Identity Percentage” of any demographic specification at any time (t) is simply:

$$IP_{(n)} = \sum_{1}^{n(t)} AP_{(n)} / n$$

30

Where $n(t)$ is the total number of queries to date involving the zip code in which the sample exists.

The following is included by way of example, without intending to limit the present invention. Below are five sample queries submitted to the present invention over 5 the first two weeks of operation, and sample results generated by the present invention. These queries involve the category "race" for the zip code 34208.

	Query #:	1
	Behavior:	watched >80% of Friends, 8:00pm – 8:30pm on 3/23/00
10	Black:	08% 16% of behavior matches
	White:	23% 46% of behavior matches
	Asian/Pacific Islander:	12% 24% of behavior matches
	Other:	07% 14% of behavior matches
	Query #:	2
15	Behavior:	muted channel at first commercial of Lakers vs Suns, 8:00 pm on 3/24/00
	Black:	02% 10% of behavior matches
	White:	10% 50% of behavior matches
20	Asian/Pacific Islander:	05% 25% of behavior matches
	Other:	03% 15% of behavior matches
	Query #:	3
	Behavior:	watched >80% of Friends, 8:00pm – 8:30pm on 3/30/00
25	Black:	27% 54% of behavior matches
	White:	03% 6% of behavior matches
	Asian/Pacific Islander:	12% 24% of behavior matches
	Other:	08% 16% of behavior matches
	Query #:	4
30	Behavior:	watched >80% of Seinfeld rerun, 7:30pm – 8:00pm on 3/31/00
	Black:	10% 17% of behavior matches
	White:	24% 40% of behavior matches
	Asian/Pacific Islander:	14% 23% of behavior matches
35	Other:	12% 20% of behavior matches

TO 5 T P D : O T I S G E R C E D O

	Query #:	5
	Behavior:	watched >40% of Seinfeld rerun, 7:30pm – 8:00pm on 3/31/00
5	Black:	22% 28% of behavior matches
	White:	11% 14% of behavior matches
	Asian/Pacific Islander:	23% 29% of behavior matches
	Other:	24% 30% of behavior matches

To further refine this example, assume that a particular STB exhibited the above behaviors as follows, and thus percentages attributed to that STB for the category of "RACE" were as follows:

Resulting *AP(n)* Contribution

15	Behavior #	n	Exhibited	BLA	WHI	API	OTH
	Behavior #1:	1	YES	16%	46%	24%	14%
	Behavior #2:	2	YES	10%	50%	25%	15%
	Behavior #3:	3	NO	06%	58%	14%	22%
	Behavior #4:	4	YES	17%	40%	23%	20%
20	Behavior #5:	5	NO	<u>23%</u>	<u>46%</u>	<u>22%</u>	<u>8%</u>
			$\Sigma AP(n)$	72%	240%	108%	79%
	<i>IP(n) =</i>		$\Sigma AP(n)/5$	14.4%	48.0%	21.6 %	15.8%

Below are the results of a sixth sample query submitted to the present invention and sample Identity Percentage (IP) calculations for the "race" category:

25	Query #:	6	
	Behavior:	watched >90% of Law & Order rerun, 7:00pm – 8:00pm on 3/31/00	
	Black:	02%	08% of behavior matches
	White:	14%	56% of behavior matches
30	Asian/Pacific Islander:	04%	16% of behavior matches
	Other:	05%	20% of behavior matches

Resulting *AP(n)* Contribution

35	Behavior #	n	Exhibited	BLA	WHI	API	OTH
	$\Sigma AP(n-1)$	1-5	N/A	72%	240%	108%	79%
	Behavior #6:	6	YES	<u>8%</u>	<u>56%</u>	<u>16%</u>	<u>20%</u>
			$\Sigma AP(n)$	80%	296%	124%	99%
	<i>IP(n) =</i>		$\Sigma AP(n) / 6$	13.3%	49.3%	20.7%	16.5%

These results would seem to suggest that STB corresponds to "White." Six queries may not be enough to decide such a conclusion. The probability that a person

actually is some IP(n) specification can be derived simply through straightforward statistical methods.

Referring again to Figure 6, the number of viewers for a set-top box may be constrained by Interval-Updating Graphic Database 620. Evolving STB Specification Percentages Database 950 can receive input from IDM Calculation Algorithm 940, which 5 may run continuously on all time-possible events.

The present invention may also utilize a second evolutionary database, illustrated in Figure 6 as Evolving STB Viewer Possibilities Database 960. A goal of Evolving STB Viewer Possibilities Database 960 is not to determine who is using a set-top box, but 10 rather who may be viewing particular content. Thus, Evolving STB Viewer Possibilities Database 960 may track possible or probable users of a particular set-top box, regardless of whether an individual, a couple, a family, or a large group of people is watching. Evolving STB Viewer Possibilities Database 960 may be updated at regular intervals, but optimally may be updated at every update of Evolving STB Specification Percentage 15 Database 950.

Data from Evolving STB Specification Percentages Database 950 may be best fit to the demographics of an area, and this best fit may be held in Interval-Updating Graphic Database 620. A best fit may be calculated by spot-filling, in which the highest set-top box specification percentages fill the most demographically significant spots first. This 20 spot-filling process continues until all spots are filled. At any time, if a category has been completely accounted, no more spots may be taken and any relevant specification can be discarded.

As a result of such calculations, Interval-Updating Graphic Database 620 ultimately holds a "best guess" at the complete specification makeup of each individual 25 in a household. Spot-filling may not be the most accurate technique, however, as a best fit reduces overall deviation. Due to higher specification percentage weighting, spot filling may provide an extremely close approximation.

The evolving database and best guess technique outlined above have been described in examples which determined viewing behaviors for a given home. However,

the present invention can also account for persons in a bar watching a football game on Sunday, even if they are already counted at home.

To determine individual persons, and not just a group percentage, matching a behavior, current set-top box viewer possibilities may be best fit to IDM calculation specification percentages for any event for the population of some graphic region or set of regions. In this way, individual behaviors can be determined. This best fit may be performed by spot-filling in a manner similar to that outlined above. Every person can be accounted for by the present invention, whether are at a bar, a neighbor's home, or their own home. The IDM specification percentage may be fit to evolutionary specification percentages for each box, thereby accounting for such deviations.

Ultimately, for each event, individuals matching a behavior may be known, and such data may be sent to Individual Behavior Determiner 970 and stored in Individual Behavior Database 290. Individual Behavior Database 290 may hold all individual behaviors recorded since sampling inception. Individual Behavior Database 290 may comprise a database of time-oriented arrays containing information about what each sample has done since sampling inception. Individual Behavior Database 290 may comprise a database approximating individual behaviors for each event from IDM Calculation Algorithm 270 or 940.

Individual Behavior Determiner 970 comprises a linear system that can find a best fit between IDM Calculation Algorithm 270 or 940 and Evolving STB Viewer Possibilities Database 960. Individual behaviors may be approximated as a best fit of the data groups from IDGM Calculation Algorithm 940 and Evolving STB Viewer Possibilities Database 960 for each event. In this manner, the behavior of one individual may be tracked over time.

New specification percentages from Evolving STB Percentages Database 950 may be continuously available and periodic recalculation of individual behaviors may be preferred. Percent changes in Evolving STB Percentages Database 950 may be small and determinable through time, and this may be the factor used to determine the interval of individual behavior recalculation.

Figure 7 is a block diagram of modules comprising Future Events Query System 190 of Figure 1. Future Events Query System 190 may be different from Past Events Query System 200. Future Events Query System 190 may comprise a web-based system, which interacts with Post Translation System 420, allowing Market Customer 530 to query the present invention regarding behaviors that are most likely to occur in the future. Future Events Query System 190 may include a web-based system with a natural language, graphical, or command-line interface, providing the customer with the ability to extract information from the system.

Individual Behavior Database 290 can contain individual-specific behavior information as determined by the present invention. Individual behaviors may be analyzed by Series Analysis System 980, which can comprise a system that looks for data trends and patterns (both directed and undirected).

Series Analysis System 980 may comprise an algorithm looking for trends in individual and group viewing behaviors that may efficiently define relevant behavior patterns. Such series analysis algorithms may be time based, defined behavioral events, and undefined behavioral events, such as a straightforward series that determines behavioral patterns. These algorithms may take individual behaviors from Individual Behavior Database 290 and look for trends based on time, content, channel changes, and the like. Series Analysis System 980 may also define relevant events according to results of its analysis.

Relevant viewing events may comprise events that represent some pattern describing viewing behaviors. Events may be pre-defined events such as "changing the channel at the end of a show" or "changing the channel at the beginning of a commercial," or they may be definitions that are more arbitrary, such as "changing the channel twice in 5 seconds before changing the channel 3 times in the next 10 seconds." Events may result from individual behavior data mining based on series analysis and behavior pattern determination, and then reporting such patterns in a simple form. By way of example, the present invention may learn that a particular graphic category may not get home until 6:00 pm, then views a news channel for half an hour, then turns the set-top box off for an average of a half hour, presumably for dinner.

One series analysis that may be performed by the present invention is a Time Series Analysis (“TSA”). In a TSA, trends that can be fully described as a function of time may be identified. This distinction is made since most mathematical series analysis methods are usually referred to as ‘time-series analyses.’

5 An alternative series analysis that may be performed by the present invention is a Defined Event Series Analysis (“DESA”). A DESA can identify trends which may be fully-described as a function of a set of behaviors, such as changing the channel near the beginning or end of an hour or near the beginning or end of certain content, watching entire programs, watching certain genres, and the like. A DESA can allow the present
10 invention to identify not only those features that are of interest to the present invention, but also to customers of the system as well.

Still another series analysis that may be performed by the present invention is a Undefined-Event Series Analysis (“UESA”). A UESA is similar in many respects to a DESA, except that a UESA can look for general trends while defining its own events. By
15 way of example, without intending to limit the present invention, sampled individuals may be on time schedules or have general viewing habits about which the present invention may soon learn.

After series analysis, events that the present invention identifies as describing behaviors in Individual Behavior Database 290 may be sent to Event Definition System
20 400, where an emphasis may be placed on behaviors that are more current. Event Definition System 400 can also accept input from Updating Future Airings Database 990 which may hold a content guide that may include content attributes and content presentation information. Such content presentation information may include, but is not limited to, to networks or channels presenting such content, and times such content was
25 or will be made available.

Event Definition System 400 can break down programming into events defined efficiently by Series Analysis System 980. Event Definition System 400 may comprise an algorithm that may accept a broad range of content attributes. Event Definition System 400 may also break content apart into a best fit of events as determined by Series
30 Analysis System 980.

Updating Future Airings Database 990 may comprise a database of arrays holding a best guess at what content may be aired at any time in the future. Airings Source 360 may continuously update Updating Future Airings Database 990. For times relatively far into the future, general extrapolations may be made to save data space.

5 Future content may be broken down in terms of viewership events existing in Event Definition System 400. From these datasets, Future Events Mapping System 450 may map future individual and group events onto future content by describing it in terms of the events defined by the Event Definition System 400.

Future Events Mapping System 450 may comprise a simple algorithm that
10 linearly forecasts the most probable events onto a map of future programming. Future Events Mapping System 450 may comprise an algorithm that takes input from Event Definition System 400 and maps probable mutually exclusive sets of behaviors which Series Analysis System 980 forecasts for these events. Events of both systems may be identically defined, thus requiring only a best fit mapping of individual behaviors onto
15 the future programming.

Market Customer 530 may query Future Events Query System 190, which may comprise a user-interface with options for tailoring such queries. Post-Translation System 420 can translate such queries into a mathematical formula that may be understood by Pre-Translation System 410. Such translations may simply express a
20 query in a format that facilitates data extraction from Future Events Mapping System 450.

Figure 8 is a block diagram of modules used in Program Entry and Program Builder Systems of the present invention. Program Entry 540 may facilitate behavioral or viewership predictions for content which has yet to be experienced by the public based on
25 data entered by a customer. Program Entry 540 may allow a customer to enter attribute ranges for certain content, and Program Entry 540 may report specific attribute values which best fit a desired outcome. Program Entry 540 may convert such customer data to a format readable by Event Definition System 400. Event Definition System 400 may break down content to determine likely viewership or other behaviors based on statistics
30 generated by other portions of the present invention.

Program Builder 430 can compile a content description, including various content attributes, for content that is likely to be popular. Such content descriptions may be based on events randomly pieced together from Event Definition System 400. Program Builder 430 may have a non-random component as well, in the form of an iterative system. An iterative system may reduce processing times and increase the likelihood of quality matches per unit time.

Event Definition System 400 can pull events that best describe individual viewing behaviors and patterns from other portions of the present invention, and such events may be entered into Random Generator 440. Random Generator 440 may comprise a random component for Program Builder 430, and Random Generator 440 may piece together content combinations to build a hypothetical program.

Random Generator 440 may comprise an algorithm which can select from content attributes and content components submitted to it in a dataset. Such selections may be performed in a computationally random manner, thereby allowing for a variety of dynamically generated content. In a preferred embodiment, Random Generator 440 may include an option, selected through Program Builder 430, which can mark as used those elements selected as part of a dataset, thereby restricting the recurrence of such elements.

As with customer-generated content, probable content popularity for content generated by Program Builder 430 may be calculated through Program Entry 540. However, this calculation may also be run in an iterative or non-iterative cycle so that an optimal "proposed" program may be described.

Figure 9 is a block diagram of modules used in a Data Mining and Prediction System of the present invention. Data Customer 510 may comprise a customer of the present invention interested in information from Prediction System 240 or Graphic Correlation System 920.

Graphic Correlation System 920 may comprise a graphic correlation database that may be updated by Series Analysis System 980. Series Analysis System 980 can analyze correlations in graphic data alone, without respect to tuner data. Graphic Correlation System 920 may hold correlations determined thus far, and may generate additional

correlations. Data Customer 510 may use graphic Correlation System 920 and its correlations.

Prediction System 240 may comprise both a system of algorithms which determines statistical probabilities of DATA1, DATA2, DATA3 or other behaviors that 5 have or will occur, and a web-based system allowing Data customer 510 to query the present invention about these behaviors or trends. Such queries may be entered through a variety of means, including natural language, graphical, or command-line interfaces. The computational algorithms of Prediction System 240 may be similar to those of Series Analysis System 980, except that Prediction System 240 may be concerned with general 10 behavior patterns, and not necessarily those having to do with television-related behaviors.

Sales Data 520, which may alternatively be seen as DATA3, may comprise sales or other operational information about business performance or trends of customers, competitors, or an industry in general. Such data may share features, such as a zip code, 15 with a DATA2 counterpart. Market Customer 530 or other data sources may provide such data in a format readable by Series Analysis System 980. This data may comprise data based on sales figures for operations in certain geographic regions, or other information, such as colors or street locations of stores, phone numbers, building types, and the like.

Market Customer 530 may comprise a customer interested in past viewership data 20 or a customer attempting to predict the desirability of previously unaired content. In a preferred embodiment, Market Customer 530 may provide additional data for Sales Data 520; such data may be cross-referenced to commercial programming, competitors, and the like.

While a preferred embodiment of the present invention is geared toward 25 measuring television viewership, the present invention may be useful for other purposes. For example, the Series Analysis algorithms used by the present invention may be run against Graphic Data, Sales Data 520, and the like without respect to set-top box data. Through such analysis, the present invention may provide detailed demographic data in 30 addition to a market research services. Such analyses may be extended to look for trends

in demographic data, thus further refining the understanding of a geographic region for advertisers, government agencies, and others interested in such data. This analysis could effectively be used to further define the effects of commercial programming, to more appropriately plan cities and city services, and other such purposes.

5 Weighting/Specification Selector 250 may determine graphic categories of interest to such parties by evaluating values from Evolving STB Specification Percentages Database 950. Low-relevancy categories may be rejected and new categories can be selected or introduced for evaluation. Weighting/Specification Selector 250 can determine graphic categories that may be most relevant to viewing behavior, and
10 may comprise an algorithm that weighs categories according to sets of specification percentage sums.

Higher mutually exclusive category sums, built on a smaller number of specification percentages, may create a higher weight. A formula may be provided that determines whether a category should be excluded in the next update of IDM Graphic
15 Matrix 910. Depending on the nature of IDM Calculation Algorithm 270 or 940, weights may be included in the calculation system rather than responsible for category exclusion in the Identity Matrix.

While the preferred embodiment and various alternative embodiments of the invention have been disclosed and described in detail herein, it will be apparent to those
20 skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope thereof.